

Multiresolution Terrain Modelling from Multisensor Data

John Wright, Kathy Sturdevant, Jack Morrison

Jet Propulsion Laboratory(JPL)
California Institute of Technology
M/S 168-414 4800 Oak Grove Drive
Pasadena, CA USA 91109-8099
FAX: (818) 393-6962,
E-mail:john.r.wright@jpl.nasa.gov

ABSTRACT

An application for registering and merging multiresolution terrain models is presented and some of the associated tools are discussed. The tools convert terrain models into voxels, store them in octrees, register the octree models to improve estimates of relative position, and then merge the models into a single virtual octree. Octrees may then be converted to triangle meshes for visualization applications. Application of the methods to extraplanetary data gathered on Mars is reviewed. In particular, the methods described are applicable to data captured by multiple sensors from multiple viewpoints with large differences in resolution and sampling rate.

Keywords: voxels, octrees, terrain models, model registration.

1.0 Introduction

Our group at the Jet Propulsion Laboratory is exploring methods of merging three-dimensional terrain datasets from a variety of sources. These datasets typically represent science data gathered on other planets and are, or need to be, used for mission planning, science analysis, and public outreach. Unfortunately, each dataset is acquired in its own coordinate space and been utilized within that coordinate space individually, with little attempt to utilize multimission data within a single application.

Future missions will require merging of these disparate datasets. Imagery collected from orbital assets, descent imagers, and planetary landers and rovers must be merged into a cohesively managed dataset. Several methods are available for this type of task but most only address part of the problem. We are trying to add the capability to handle the types of data that existing methods do not currently handle well. For example, the

datasets collected from the orbital assets have a long history of use and many tools have been developed for working with them [LAVOIE]. These datasets are generally classified as 2.5D or gridded models with elevation and multispectral (e.g. RGB, IR, UV) data being stored as images. Many standard image processing techniques are available to register these images, perform stereo elevation extraction, perform orthorectification, etc. Descent imagery falls into the same category in that it is generally collected while looking straight down or only slightly obliquely and many of the same tools can be applied. However, imagery taken from a lander or rover is very different. Because of the proximity to the surface, and the oblique angle of the image, standard image processing techniques fail at the registration process. At this point we begin to explore registration techniques that work with three-dimensional models.

Our method uses voxels to represent the terrain and our registration and merging operations are done in the voxel realm. We selected voxels because of certain basic advantages they offered over other methods. Many of these advantages, and the application of voxels to model registration, were outlined in [CURLESS] and [HILTON97]. The primary advantage is the ability to encode voxels as representing known terrain surface, known empty volume, or unknown volume. This allows the registration of terrain models which were generated from different viewpoints and have areas which are only visible from one location. For example, when the lander captured an image of the rock Yogi during the Pathfinder mission, the volume behind the rock is unknown. If the rover were to drive partway around the rock and capture an image of the back side, the volume would be known to be either empty or surface. Having empty and unknown volumes, in addition to the known surface, allows us to perform a better matching process. A second advantage of voxels is the relatively straightforward conversion from existing datasets into voxel formats. For gridded elevation data, each point on the grid is converted to a voxel whose edge length is equal to the pixel spacing. For the lander and rover stereo imagery, the range maps can be used to compute the volume of the pixel view frustum at the range to the terrain. This volume readily becomes a voxel for our dataset. The empty volumes are generated as all the volume above the gridded elevation data or the volume of the view frustum from the viewpoint to the terrain in a range map.

Because voxels generated from a range map will range across a variety of sizes based on range from the viewpoint, it was considered necessary to use a voxel structure which supported multiresolution data. This, and the fact that the known data is generally a thin set of surface voxels and that large parts of the volume will be uniformly empty or unknown, immediately suggested the use of octrees for encoding the voxels in the volume. This provides a relatively compact form for storing the voxel data, inherently supports multiresolution, and allows easy merging of multiple datasets.

2.0 Motivation

The Jet Propulsion Laboratory is responsible for the operation of many extraplanetary exploration missions. During the Pathfinder mission to Mars it became apparent that existing methods for working with terrain data for mission planning were not effective. In particular, it was extremely difficult to merge terrain data from past missions, such as Viking and Voyager, with data captured by the lander and rover cameras. Viking data has been used to map the entire planet at about 450 meter resolution and establish latitude/longitude conventions. When Pathfinder landed, its position was estimated, using visible landmarks, to an accuracy of about several kilometers. Unfortunately, the resolution of the three-dimensional terrain models constructed from data from the onboard stereo cameras is much higher than the resolution of the existing terrain data or the positional error, on the order of centimeters. Thus, it was effectively impossible to merge the terrain models and claim that the Pathfinder data contributed to a higher resolution model of the Martian surface. All modelling of the terrain around the lander was done in lander-centric coordinates, with only an estimate of North to tie the models to Mars.

Efforts are underway to improve this process since better terrain knowledge will be necessary for future mission operations. Pathfinder was one of the first missions to combine autonomy for onboard navigation decisions with overnight processing and analysis of data gathered for planning the next day's activities. As the levels of autonomy increase, the range of the rover will increase and larger terrain areas will be explored. This will take the rover outside of the three-dimensional terrain model built from the lander's stereo camera data and it will be essential that some knowledge of that terrain be available for use in planning activities. Also, the increased range will increase the amounts of science data

being returned and increase the overnight processing workload. Pathfinder was fortunate in that the original terrain model was built with the first suite of images captured by the lander's cameras. The remainder of the image data collected during the mission was used to enhance the original terrain model but the operators always had at least a rough model of any region to be explored. This will not be the case in longer range missions as the rover may move completely out of known regions.

As the range of the rovers increases, it will be necessary to utilize additional data to have some information to use in planning activities. The existing Viking data is unlikely to be very useful as the resolution is still much lower than desirable. However, Mars Global Surveyor is mapping the planet at much higher resolution, on the order of 100 meters, and is even collecting imagery of some areas at submeter resolution. This imagery will be very useful, provided it can be registered with the Viking data and that the lander can be properly located within it. Another type of data, which was not available during the Pathfinder mission, was descent imagery. As the lander descends to the surface it takes a sequence of images which are transmitted back to Earth. Efforts are currently underway at JPL to process a sequence of descent images to generate elevation models for the terrain being imaged. As the lander descends, each image covers a smaller area so the resolution improves but the area covered at that higher resolution is reduced. Thus the mission operators will have fairly high resolution data immediately around the landing site which gradually decreases as the rover ranges farther afield.

Of course, the stereo cameras on the lander and/or rover will also contribute data for constructing terrain models for activity planning purposes. The rover will drive to a new area, collect a panorama of images, and transmit the images back to Earth for processing. This data will be converted to a terrain model and local science activities can then be planned and accomplished. However, if the rover is to collect specimens and return them to the lander for launch into orbit, the rover's (and lander's) position must be known fairly accurately. Pathfinder used internal gyros and inclinometers to perform dead reckoning but the operators had daily imagery from the lander to use to correct any errors. When the rover wanders farther afield, this will no longer be available. Since the Martian GPS system will not be in place for some time, it will be necessary to use knowledge of the terrain for navigation.

Some work is underway to process the descent imagery to create a set of landmarks to be used for navigation. This will be extremely useful but methods of tying the landmark database to the local terrain models will be needed.

In general, all these datasets need to be integrated such that each dataset is georeferenced and the mission planning tools can use all available datasets for defining upcoming activities. It will no longer be possible to work within a single isolated coordinate space given the rover's wider range and the varied sources of data available. The task whose work is described here is intended to merge these varied datasets and provide geolocations to use as the common coordinate space.

3.0 Implementation

The System for Unifying Multiresolution Models and Integrating Three-dimensional Terrains (SUMMITT) task has the goal of developing the underlying modelling technology for supporting missions involving rovers. Three-dimensional models of terrain areas are an invaluable asset for planning operations and reviewing the predicted and telemetered operations of a robot arm. The SUMMITT task had the initial goal of supporting the Mars Volatiles and Climate Surveyor (MVACS) team during the Mars '98 mission, which unfortunately failed. The next Mars surface operations missions are the Mars '03 missions that plan to land two rovers in January and February of 2004. These missions expect to have orbital imagery from Mars Global Surveyor (MGS) and Mars Odyssey as well as lander imagery from a stereo imager similar to the Imager for Mars Pathfinder (IMP) used during the Pathfinder mission. These sources of imagery will be combined to create a multiresolution terrain model with very high resolution detail available within the immediate area of operations of the rover.

Figure 1 gives an overall block diagram of the terrain modelling system. The three primary sources of data are orbiter imagery, descent imagery, and lander imagery. Each type of imagery is partially processed independently, then combined with data derived from the other imagery to create the multiresolution terrain models. Many of the processing steps described in the block diagram and in the following text have been developed as part of a variety of tasks within JPL. A key goal of this task is to identify the portions already developed and establish the

necessary tools to integrate those technologies into the overall system.

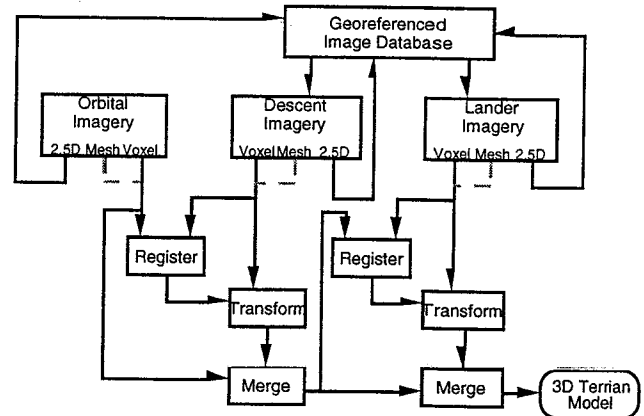


Figure 1: Terrain Modelling System Block Diagram

3.1 Creating Voxel Models

There are two primary sources of data used for terrain modelling: 2.5D gridded elevation data and range data. Gridded elevation data is typically generated from stereo imagery captured by orbiting platforms or other high altitude imagers. The stereo data is then processed into range imagery that is then converted to a grid of elevation values. The gridded elevation data is typically orthorectified and is often georeferenced (tied to the latitude/longitude grid of the planet). However, other sources of data can also be used to generate gridded elevation data. These include laser or radar altimeters (e.g. the Mars Orbiting Laser Altimeter) or synthetic aperture radar (e.g. Magellan mission to Venus).

Range images are typically captured by stereo imagers on the planetary surface, either on a stationary lander such as Pathfinder or mounted on a rover. The stereo imagery is processed to compute the range to each pixel in the image. Because the difference between the minimum and maximum range values across the image is typically relatively large, as compared to data captured from orbit, and the images are typically captured at oblique angles, these datasets are stored as range images rather than being converted to a gridded elevation format.

Conversion of these models into a voxel form is relatively straightforward. For gridded elevation data, each point in the grid becomes a voxel with a center at the (x,y,z) location of the grid point and an edge length equal to the distance between an adjacent pair of grid points. For a range image, each image pixel is projected along the pixel ray

from the camera position to the distance specified by the range value. The (x,y,z) position is computed based on imager position and orientation and the voxel edge length is computed as the pixel view frustum width at that distance. This process is more complicated than that for the gridded data because precise models of the imager are required to properly compute the projected point locations.

As each voxel is generated, it is inserted into an octree. The octree is sized such that the voxels occupy most of the octants and the voxel (x,y,z) positions are transformed into octree space prior to insertion.

The octree acts as a data structure for storing the voxels rather than as a space subdivision structure. In essence, it acts as a set of bins that hold the voxels and provide rapid search and access mechanisms. Each voxel inserted into the octree retains its original (x,y,z) position and is stored in the subcube of the octree that contains that (x,y,z) position. Each subcube can contain multiple voxels. The edge length of the voxel determines the level to which the voxel descends before being added to the subcube. Octrees provide significant advantages for storing 3D terrain data. Octrees provide an inherent multiresolution data structure that supports the concept of merging multiresolution data. Octrees also provide relatively compact representations for sparse data. The most important part of the terrain models, for the applications being considered, is the surface. It is unnecessary to store more than that for many applications. However, it can be quite useful to also store empty and/or unknown volumes in the octree to aid in registration.

Figure 2 shows a simple example of an early voxel model for a section of air bag and terrain patch as seen from the Pathfinder lander imager. The individual voxels are reduced in size to make it easier to see the nature of the data. Figure 3 shows a more sophisticated example generated from a perspective view and illustrating how hidden surfaces show up as missing areas in the model.

Typical applications of these terrain models will be visualization for mission planning and rehearsal, science data analysis, and public outreach. The principal tools for these applications rely on polygon models for high performance rendering. Section 3.3 describes the work on converting octree models to polygon models. However, to save effort and computations, the pixel connectedness of the original image data may be used to generate a polygon mesh for each individual dataset. This adds about a 10% overhead to the octree generation process but is expected to save much




Figure 2 - Simple Voxel Model

more time during the mesh generation phase later. Thus, the final model generated during this step is a combination of an octree and a triangle mesh which we call a surface model.




Figure 3 - Perspective Voxel Model

3.2 Registration

The fundamental problem is the registration of the different terrain models generated independently from the different data sources. For essentially orthonormal views, such as orbital and descent imagery, simple image correlation works well and has been extensively utilized. The orbital data becomes the baseline 3D model and the descent data is used to enhance the baseline. However, processing the lander/rover imagery is more difficult. Two techniques were considered for performing this process and registering the models from the lander imagery to the baseline model. One technique is to reproject the models into an orthonormal view and then use two-dimensional methods to register the images to orbital and descent imagery. A difficulty is the inherent change in resolution across the orthonormalized image and research is needed to resolve this problem. Another significant problem for image correlation methods is the effect of lighting variations. Differences in shadows and shading due to changes in sun position can cause problems for image correlation processes.

A method that has been demonstrated and selected at JPL is 3D registration of the 3D models. Registration in 3D space offers significant advantages over 2D registration in image space. Typically, the stereo data for generating 3D models is acquired over short periods of time or simultaneously. Thus, the lighting conditions are relatively constant and 3D model generation is fairly good. The 3D models may then be registered with no concern for lighting variations. Also, 3D shape registration can be done for models with more than an order of magnitude difference in resolution. Methods for registering polygonal surface models have been explored. These techniques use optimization techniques to match the 3D surface of the terrain model generated from the lander imagery to a similar 3D surface generated from the baseline model. The resulting match parameters are then used to correct the position and orientation information for the lander. The selected approach uses volumetric primitives (voxels) to represent the terrain to be matched. Voxels have some advantages over the polygonal surface matching methods in that it is easier to represent unknown volumes, such as regions occluded by rocks and hills, and easy to use, multiresolution data structures are available in which to combine the models once they are matched. These methods are still under development but have been proven to be effective. The images in figure 4 display the results of registering a baseline model, with coarse, large voxels, with a higher resolution terrain model generated from a low-level, oblique view.




Figure 4 - Voxel Model Registration

The view on the top shows the two models overlaid with an artificially induced error in the camera position. The view on the bottom shows the result of the registration process and the accuracy of the alignment. The registration uses an iterative closest points method based on the work of [BESL], [ZHANG], and [CHAMPLEBOUX]. The voxels from the oblique view are colored red to facilitate distinguishing them from the baseline voxels but appear as darker in grayscale images.

herein. Note that the red/darker voxels are significantly smaller and more dense than those of the baseline model. The entire voxel model is stored in an octree structure which supports multiresolution data and rapid access.

Once a model has been registered to the baseline model, the combination of the models is stored in a simple list structure. Each entry in the list represents a surface model and the transform needed to align it with the baseline model (the first entry in the list). This collection of octrees is called a forest. It is a virtual merging of the octrees into a single octree but supports the capability of adjusting the alignment transformations based on additional registration information acquired later. The forest then becomes the baseline model for the next registration step.

The preceding discussion primarily concerned lander imagery but is equally applicable to rover imagery. On Sojourner, the stereo imager was very small and no attempts have been made to generate 3D terrain models from its imagery. This was done exclusively with the IMP imagery from the lander. For Mars '03, the rovers are expected to carry IMP-like imagers to capture high-resolution stereo imagery of the terrains under exploration. In these cases there will likely be two modelling paradigms. In one, the rover will traverse to a new region, then stop and capture a panoramic image for creating a local model. Identification of landmarks previously noted by the lander and rover will be used to triangulate the rover's position. If the area is completely new, terrain matching will be performed between the high resolution model, generated from the stereo imagery, and the baseline model. However, if the new region partially overlaps a previously explored region, then terrain matching techniques may be used to register the overlapping sections and thus identify the new position of the rover. In a more localized case, the rover may move to image a region occluded by a rock at the previous position. In this case, a large portion of the 3D models generated from each location will overlap, with only the occluded portion being different. Terrain matching will work well in this case to integrate the newly revealed and modelled area into the baseline model.

3.3 Application of Models

Once a terrain model has been generated, it is typically converted to a polygon model for utilization by a visualization tool. Because the operator will be making planning decisions that require detailed local knowledge combined with general understanding of more distant terrain, the

polygon models must be multiresolution also, or at least reflect the multiresolution nature of the underlying data samples. The multiresolution nature of the data precludes the use of a simple algorithm such as Marching Cubes [LORENSEN]. To extract the polygon model from the octree, a modified version of the Marching Triangles algorithm [HILTON96] was developed. The modified version utilizes only points in the model to generate a Delaunay triangulation in three-space. By using points in the model rather than selecting points on the implicit surface described by the model, the multiresolution nature of the data is preserved and reflected in the triangulation. Figure 5 shows an example of a multiresolution model converted to a multiresolution triangle mesh, both shaded and wireframe. From the wireframe model it is possible to see that the oblique view of the hill was from the upper right as the small triangles are on that side of the peak.

One concern about the Marching Triangles algorithm is the compute time required to generate a triangle model from a large number of points. The algorithm uses a search of previously processed points that is $O(n^2)$ so large models require a large amount of time. Use of hash tables can reduce the total time but not the order of the problem. One way around this limitation is the use of the pre-generated meshes created along with the octrees. These meshes are generated from the original image pixel connectedness and thus reflect a "good" triangulation of each piece of the data. Thus, the problem may be reduced from creating a large mesh to stitching together a number of smaller meshes. Figure 6 shows a closeup view of a boundary between a low resolution mesh and a high resolution mesh stitched together with tools developed for the task.

Mesh reduction adds another aspect to the problem. If every point in the original datasets becomes a vertex in the output mesh, the number of triangles becomes very large. Mesh reduction is necessary to keep the meshes to a reasonable size. Octrees provide a handy way to perform mesh reduction. Simply averaging the voxels contained within each block of eight subcubes in the octree and pushing the resulting average voxel up the tree to the next level provides a simple way to reduce the resolution of the dataset. At each step in the accumulation process, a mesh can be generated providing a multiresolution mesh for high performance rendering. The method is partly adaptive in that it simplifies the high resolution portions first. It also can be performed only on the lowest levels of the octree to provide some noise reduction through averaging. Table 1 illustrates the resulting reduction in voxels and

triangles for a series of accumulation steps on a multiresolution model.

Accum. Steps	Number Voxels	Number Triangles
1	90665	178389
2	39885	79109
3	18184	35839
4	4751	9242
5	1243	2353
6	326	584

Table 1 - Simplification by Voxel Accumulation

Another aspect of mesh reduction is texture mapping. When generating the original octrees and surface models, the pixel color may be inserted into the voxel for later display. However, when the mesh is reduced, it is often desirable to use the originally captured images as textures for the mesh. The simple case is to use a single texture image created as a combination of some set of the original images. However, this does not readily support multiresolution applications. More effective is the use of multiple texture images mapped onto the corresponding portions of the mesh. Figure 7 illustrates an example of using three images as texture for three registered surface models.

An additional mesh generation technique has been developed for processing the terrain models. This method is used to generate Active Surface Definition (ASD) models for use by Performer™, a high performance visualization library developed by Silicon Graphics, Inc. The Performer library is being used to develop tools for visualizing the surface of Mars and creating command sequences for directing rover activities. Due to the nature of the mission and the requirements for success, it is desired that the visualization tools have high fidelity and high rendering performance. ASD models help provide this capability by utilizing multiple Levels of Detail (LOD) within the model and rendering only the appropriate detail for the scene and range from the eyepoint.

Typically, tools for generating such models use subdivision surface methods and the approach used here is no exception. It is typically simpler to start with a very primitive triangle mesh and add points to it to get a progressively more complex mesh. Many methods, such as that described by [SCHRÖDER], require a regular mesh. To generate such a mesh, the initial irregular mesh is simplified to a small set of triangles and then each triangle becomes a patch that is regularly sampled back to the original mesh. These methods have the benefit of preserving the topology and genus of the original model. However, they do not necessarily utilize the original set of mesh points and thus




Figure 5 - Multiresolution Mesh Example




Figure 6 - Stitched Mesh Boundary

may not quite match the original model. For our purposes, it is desirable to utilize the original mesh points as these are known data points collected by our sensors and the overlap or integration characteristics of the sensor elements may be unknown.

To accomplish this, the ASD model generation tool assumes a planar topology for the terrain surface and begins with two triangles. Each triangle is subdivided by picking the original mesh point closest to the center of the triangle and generating three subtriangles using the three original vertices and the new point in the center. This step repeats until the triangulation reaches the desired resolution or no additional subdivisions occur. Constraints on the selection of the center point prevent the generation of "bad" triangles that make the surface too noisy. Rapid selection of the center point is supported by the octree structure, providing $O(\log n)$ search capability.

Figure 8 illustrates the results of the ASD surface model generation process with multi-resolution texture applied to the resulting triangles. There are three levels of mesh detail visible in the image, with the highest resolution in the foreground, and two levels of texture resolution visible.

These examples have illustrated the use of octree and surface models to generate triangle meshes for visualization applications. Other applications require gridded elevation datasets for visualization and analysis purposes. Generating gridded elevation datasets can be done by sampling the octree volume bottom up and inserting the data values into the elevation image. This step is followed by an interpolation/hole filling step to complete the image.

These techniques are currently used or are under development at JPL and are planned to be in place to support the upcoming Mars missions beginning with Mars '03.

4.0 Results

The tools developed for working with terrain models in octree and mesh form and converting the models to meshes for visualization have undergone extensive development and refinement. Performance improvements have increased the execution speed for most of the individual components. For example, the early prototype required approximately twenty minutes to generate the mesh depicted in Figure 5 and containing about 2000 voxels. The current version can produce octree surface models from one gridded elevation model and two range images, coregister them, and generate a multi-resolution mesh (as depicted in Figure 7) in under ten minutes and processing a total of about 100,000 voxels. However, more performance improvement will be needed to support upcoming missions as described in section 6.0.

5.0 Conclusion

The Jet Propulsion Laboratory is moving forward with efforts to develop tools for creating models of the operational environment and providing visualization tools to explore and interact with that environment. Multi-resolution terrain models are required to provide rover operators and mission scientists with the knowledge necessary to plan science and exploration activities with minimal mission risk. The use of voxels and octrees provides a new capability to merge terrain data from multiple missions and multiple instruments into a single, unified, coregistered dataset. Such datasets may then be converted to other forms, such as triangle meshes, for high performance visualization.

6.0 Future Work

There are many areas to be explored and developed to continue this work. One key aspect is the use of real world data. The sample data used here contains very little noise and uncertainty. There are strict error bounds on camera position and orientation. Some real world data has been collected and processed and it shows considerable noise and other problems in the image data. In addition, rather than using a stereo imager with a fixed camera separation and well known camera model, this data was collected by moving the imager between frames. This paradigm is likely during rover operations to



Figure 7 - Multiresolution Texture Example

collect range information on more distant objects and landmarks. However, this introduces larger errors into the stereo separation and thus the range computations. Differences in scale between models of the same area become apparent and must be accounted for in the registration process. This process currently only models rotation and translation effects.

Another issue to be explored is the use of other meshing algorithms. Marching Triangles works fairly well but the initial triangle selection is key and needs to be automated. As the octree really represents a cloud of points, other meshing algorithms should be applicable and need to be examined, such as that of [HOPPE] or [KOBELT].

A key aspect for actually supporting the Mars '03 rover mission is the volume of data. For example, the IMP was a 256 x 256 pixel imager with about a 15° field of view. Thus about 72 images were required to generate a single

mosaicked panoramic image. For Mars '03, the main imager will be 1024 x 1024 pixels and it will have a field of view of about 16° with more overlap than Pathfinder thus requiring about 90 images to generate a panorama. This means that the data volume will be about 20 times as much or about 90M voxels for a given local operational area. In addition, there will be several stereo imagers on the rover. There will also be imagery captured by Mars Global Surveyor at about 1 meter resolution. This data will need to cover an area about 10 km on a side to cover the possible sortie range of the rover thus adding about 100 million voxels to the model. Descent imagery is not expected for this mission but in future missions it will push the voxel requirements even higher. Thus methods are needed to allow operations on files rather than attempting to load the entire model into memory. In addition, more rapid access and lookup mechanisms will be needed, primarily in the mesh generation area.




Figure 8 - Multiresolution ASD Model with Multiresolution Texture

ACKNOWLEDGEMENTS

The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- Besl, P.J. and McKay, N.D. (1992) A Method for Registration of 3-D Shapes, IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 14, No. 2, pp 239-256. February 1992.
- Cooper, B. (1998) Driving on the Surface of Mars Using the Rover Control Workstation, Proceedings of SpaceOps '98, Tokyo, Japan, June, 1998.
- Champleboux, G., Lavallee, S., Szeliski, R., and Brunie, L., (1992) From Accurate Range Imaging Sensor Calibration to Accurate Model-Based 3-D Object Localization, Proceedings of IEEE CVPR '92, pp. 83-89, 1992.
- Curless, B., and Levoy, M., (1996) A Volumetric Method for Building Complex Models from Range Images, Proceedings of SIGGRAPH '96, pp. 303-312, 1996.
- Hilton, A., and Illingworth, J., (1997) Multi-Resolution Geometric Fusion. Proceedings of International Conference on Image Processing, 1997.
- Hilton, A., Stoddart, A.J., Illingworth, J., and Windeatt, T. (1992) Surface Reconstruction from Unorganized Points. Proceedings of SIGGRAPH '96, pp 71-78, July 1992.
- Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., and Stuetzle, W. (1996) Marching Triangles: Range Image Fusion for Complex Object

Modelling. Proceedings of International Conference on Image Processing, 1996.

Kobbelt, L., Vorsatz, J., and Seidel, H.-P., Multiresolution Hierarchies on Unstructured Triangle Meshes, Computational Geometry Theory and Applications,

LaVoie, S., et al, (1999) Processing and Analysis of Mars Pathfinder Science Data at the Jet Propulsion Laboratory's Science Data Processing Systems Section. Journal of Geophysical Research, vol. 104, no. E4, pp. 8831-8852, April 25, 1999.

Lorensen, W. and Cline, H. (1987) Marching Cubes: A High Resolution 3D Surface Construction Algorithm. Proceedings of SIGGRAPH '87, pp. 163-169, July 1987.

Schröder, P. and Zorin, D., Eds. Course Notes: Subdivision for Modeling and Animation, ACM SIGGRAPH, 1998.

Wright, J., Hartman, F., and Cooper, B. (1998) Immersive Environments for Mission Operations: Beyond Mars Pathfinder, Proceedings of SpaceOps '98, Tokyo, Japan, 1998.

Zhang, Z. (1994) Iterative Point Matching for Registration of Free-Form Curves and Surfaces, International Journal of Computer Vision, Vol. 13, No. 2, pp 119-152, 1994.